

# Combining Ontology Queries with Key Word Search in the GloServ Service Discovery System

Knarig Arabshian Computer Science  
Department  
Columbia University  
New York, NY 10027  
knarig@cs.columbia.edu

Henning Schulzrinne Computer Science  
Department  
Columbia University  
New York, NY 10027  
hgs@cs.columbia.edu

## ABSTRACT

GloServ is a global service discovery system which aggregates different types of services in a globally distributed network. It improves on current service discovery systems by scaling across a globally distributed network and allowing intelligent querying and registration of services. It uses the Web Ontology Language Description Logic (OWL DL) to classify services in an ontology and map knowledge obtained by the ontology onto a scalable hierarchical peer-to-peer network. We present an enhanced novel querying mechanism for service discovery which combines ontology queries with key word search. The ontology-based query is transformed to a first order predicate logic query which is used for routing the query to the appropriate servers. Matching instances are further filtered and ranked by matching of the user's key word queries to key words in the service instances. Currently, querying is limited to either simple attribute-value pair searches, structured ontology queries or key word search. Combining ontology queries with key word search allows structured and unstructured data to be queried for simultaneously which provides users a way to query for a service in greater detail and enhances their experience when searching for services.

## Keywords

service discovery, ontologies, OWL, CAN, peer-to-peer, key word search, ontology queries

## 1. INTRODUCTION

Current service discovery systems use simple attribute-value pair matching in order to discover services, which limits the results only to exact matches. They also do not scale well and are limited to local area networks. However, as more services become available, service discovery in a wider area network is necessary and network scaling becomes an issue. The proliferation of services also creates the problem of finding different relationships between services and performing intelligent query matching. Services may also be dynamic in nature which requires the system to handle frequent service updates. In addition, context-aware service discovery in ubiq-

uitous and pervasive computing systems require wide availability of services as well as the ability to perform reasoning based on a user's context.

In order to address these problems, we have developed GloServ [9], a global service discovery system, which uses the Web Ontology Language Description Logic (OWL DL) [3] to classify services in an ontology and map knowledge obtained by the ontology onto a hybrid hierarchical peer-to-peer network. It operates in wide as well as local area networks and supports a large range of services that are aggregated and classified in ontologies. A partial list of these services include: events-based, physical location-based, communication, e-commerce or web services. Organizing services in an ontology and searching within that ontology allows searching for general categories of services and then specializing to specific services.

Querying within GloServ is performed by initially routing an ontology-based first order predicate logic query to the servers which handle that service class. Each server holds an ontology for a particular service class which has a number of subclasses restricted on various properties. For example, a Restaurant service class has a list of subclasses restricted by location, cuisine or both. When the ontology query reaches the correct servers, it is represented as a temporary class within the ontology and classified using a reasoning engine. The reasoner outputs classes that are related to the query. Relationship between classes is defined by semantic relationships such as subsumption, subclass, or equivalence. Instances which belong to these classes are returned to the user. In addition, the ontology query gives the user the option to search for similar matches such as instances belonging to sibling classes to the query or disjoint matches which are instances belonging to classes that are disjoint from the query class.

GloServ also performs key word search in addition to ontology querying. Key word searching further refines the results by querying for certain terms within a service instance's property. We provide a way of performing key word search on the ontology constructs themselves by allowing service providers to add key words to classes within the ontology.

The main contribution of this paper is allowing the combination of ontology queries with key word search. Currently, querying is limited to either simple attribute-value pair searches, structured ontology queries or key word and text searches. Combining ontology and key word queries gives more flexibility and accuracy when obtaining query results. Performing pure ontology queries limits a service description and query to the the service ontology definition, but it gives inferencing capabilities which allow for more sophisticated queries. On the other hand, when a service is described purely with text, although all types of descriptions can be added, an obvious drawback is that the correctness of the results can only be

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Middleware 07, November 2630, 2007, Newport Beach, California, USA.  
Copyright 2007 ACM 978-1-59593-935-7 ...\$5.00.

approximated and may or may not be what the user is looking for. Thus, when combining structured and unstructured querying, we reap the benefits of both mechanisms by first performing an ontology query in order to route the query to the correct servers, honing in on the correct set of service instances and then further refining this set with key word matching on a set of key words.

Below, we describe the combined ontology and key word matching mechanism. Section 2 gives background information on ontologies and the CAN peer-to-peer network. Section 3 motivates the underlying reasons for building an ontology-based global service discovery architecture and Section 4 gives an overview of GloServ. Sections 5, 6 and 7 describe the ontology querying, key word matching and implementation respectively. Related work is discussed in Section 8. Finally, we conclude in Section 9.

## 2. BACKGROUND

### 2.1 Ontologies

An ontology is defined as a formal, explicit specification of a shared conceptualization [13]. A conceptualization refers to an abstract model of how people think about things in the world. The concepts and relationships are given explicit names and definitions. These are formalized into a specification which is encoded in a logic-based language. Ontologies are meant to be shared across different applications and communities.

An ontology specification includes a number of classes which represent various concepts, similar to how object orientation or the entity-relation model consider classes. Classes have object or datatype properties which can be restricted using existential or universal quantifiers. Classes contain instances or individuals which represent actual data. Description logic ontologies, such as OWL DL, allow class relationships to be inferred based on established relationships within the ontology. Relationships can also be established between classes via logical connectives such as intersection, union, or complement.

### 2.2 Content Addressable Network

The GloServ architecture is designed as a hierarchical peer-to-peer architecture. We use a Content Addressable Network (CAN) [22] for the peer-to-peer network. A CAN is a fault-tolerant, scalable and self-organizing distributed peer-to-peer network, formed as a  $d$ -dimensional torus separated into a certain number of zones. The coordinate space is dynamically partitioned among all the peers such that every peer possesses its individual, distinct zone within the overall space. Each peer in the CAN maintains a routing table that holds the IP address and virtual coordinate zone of its neighbors. When a query comes into a CAN node in the form of a dimension-key value, the node checks to see if it holds information for that dimension-key pair. If it does not, it routes the query to the neighboring node which is closest to the destination coordinates. The worst case number of message hops in a CAN is from one end of the torus space to the other, which evaluates to  $O(n^{1/d})$  hops. If the number of dimensions is fixed to be  $(\log_2 n)/2$  then the run-time reduces to  $O(\log_2 n)$ . We have employed this in our algorithm to assure scalable routing.

## 3. MOTIVATION

The motivation behind using ontologies for service discovery rather than using simple attribute-value representations of data, such as in traditional databases is mainly due to the reasoning power behind ontologies. An example of a query which can be done using an ontology which is difficult to do using an SQL query would be

something like: “Given a service class, find all related matches to my query”. Service search is evolving to cater to particular users and their context and thus reasoning capabilities such as these are essential for performing context-aware searches. SQL also does not support abstract data types, thus making it difficult to determine if a certain property value belongs in a number of different classes or types. Ontologies can also be shared, re-used and changed flexibly. For example, when new relationships are established within the ontology because of ontology migration or the addition of new classes, determining new relationships within the ontology simply reduces to running a reasoner on the ontology in order to reclassify the classes.

The main drawback to using an ontology is that classification is expensive. As ontologies grow large and especially when instances of classes are stored in the ontology, reasoning becomes a bottleneck. We tackle this problem by storing instances in a database back-end instead of the ontology itself and using only class relationships for determining which classes a query belongs in. This speeds up the classification process considerably. Also, a positive side-effect of the distributed architecture of GloServ allows each server to handle an ontology for one service class. Thus the size of the ontology always remains manageable for classification.

A number of benchmarking tests have been performed on the currently available description logic reasoners. [10] records performances of current description logic reasoners on the top 10 most difficult ontologies. The one that matches closest to our model is a description logic ontology which has around 1500 classes and 0 instances. Such an ontology is classified in about 1 second with the Racer-Pro reasoning engine [15]. We anticipate the number of classes for a given service class to be much less than 1500 classes. For example for location-based services such as restaurants or theatres, the services may be distributed by location and then by cuisine, or movie genre. Given that there are around 600 cities [1] with population greater than 50,000 in the United States and around 1300 urban areas in the world [2], even if all the locations are put into one ontology, the number of classes will still be manageable. However, this is highly unlikely given that services are distributed in GloServ and each subnetwork handles a subset of a high-level service class. Thus, we anticipate an average service ontology to have a few hundred classes which will classify very quickly. Furthermore, below we describe an algorithm which converts an ontology query to numeric CAN keys thus requiring query classification to only occur once within a CAN network.

The motivation behind designing a hybrid hierarchical peer-to-peer architecture is to provide an efficient architecture for global distribution of services that may also be dynamic in nature. We envision all types of services to be handled within GloServ and these include allowing updates of service descriptions in real-time. For example, restaurants may want to update their available seating every 15 minutes during peak hours. In New York City alone, there are 8,000 dining establishments which results in 32,000 updates per hour for only restaurant services. If data is replicated across servers, large number of updates of all types of services will not scale.

Furthermore, we describe a novel mapping algorithm in [9] that combines the benefits of OWL DL and CAN to map content of service instances to nodes in a peer-to-peer network. Although there are other types of structured peer-to-peer networks such as Pastry [24] and Chord [25], we have elected to use CAN because it is easily constructed given a service ontology. Organizing similar classes within a peer-to-peer network causes them to lie in closer proximity to each other which makes the similarity searches defined above faster. Thus, in order to achieve load distribution, fast

query and update processing time, while maintaining reliability, we have elected to use a hierarchical peer-to-peer network as the underlying service discovery architecture.

## 4. OVERVIEW OF GLOSERV

The GloServ service discovery system achieves large-scale distribution of semantic service data that is queried for with specificity and efficiency, due to its hybrid hierarchical peer-to-peer architecture and ontology service descriptions. These attributes make GloServ a very good candidate for context-aware applications. Such context-aware applications need to have access to both ubiquitous as well as pervasive information, thus revealing the need for a globally scalable system. We give an overview of the design in this section but encourage the reader to refer to [9] and [8] for greater detail as the remaining sections concentrate on the querying enhancements made to GloServ.

### 4.1 Design

There are many ways to engineer ontologies. We have adopted the modularization approach specified in [16] and [23]. Modularizing ontologies into separate domains allows ontologies to be re-used, maintained and to evolve with flexibility. Modularization is achieved by putting general classes within an ontology in a pure hierarchy where siblings are disjoint from each other. This creates a *primitive skeleton*. Hence, service instances will only be classified within one of the branches. At the lower levels of the ontology, classes may have relationships with other classes and a pure hierarchy is not maintained. The upper hierarchical ontology which defines high-level services is mapped onto a hierarchical network and the low-level ontologies are mapped to a peer-to-peer network.

A GloServ server (GloServer) in the hierarchical network represents a high-level service class within the primitive skeleton the service ontology, which is the pure hierarchy described above. Thus, it works similar to DNS except that the network exploits the knowledge obtained by the service classification ontology to create the hierarchy. Also, since the high-level service classes are disjoint, a query will be routed to one of the hierarchical branches, largely reducing the number of hops a query needs to propagate through.

The lower levels of the network architecture is organized in a peer-to-peer network and also represents the class relationships within the ontology. We have elected to use CAN because it distributes data according to content and is easily constructed given a service ontology. Organizing similar classes within a peer-to-peer network causes them to lie in closer proximity to each other which makes the similarity searches defined above faster.

GloServers maintain three types of information: a service classification ontology, a thesaurus ontology and, if part of a peer-to-peer network, a CAN lookup table. The high-level service classification ontology is not prone to frequent changes and thus can be distributed and cached across the GloServ hierarchical network. Each high-level service will have a set of properties that are inherited by all of its children. As the subclasses are constructed, the properties become specific to the particular service type. The thesaurus ontology maps synonymous terms of each service to the actual service term within the system.

Services are represented as instances of the service classes and usually reside in the more specific, lower levels of the ontology. Each service instance has a set of properties that are populated. According to the service's attributes, it is classified in a set of related classes within the ontology. Services are registered either in a user-centric way through a web-based form or in an automated fashion by issuing a first-order predicate logic query.

At the lower levels, maintaining a purely hierarchical ontology

structure becomes difficult as classes tend to overlap. Thus, in order to efficiently distribute service instances according to similar content, servers that hold information on similar classes are distributed in a peer-to-peer network. We employ a CAN peer-to-peer architecture to distribute classes with similar content. The CAN architecture is generated as a network of  $n$ -level overlays, where  $n$  is the number of subclasses nested within the main class. An example of an ontology classification using the *Restaurant* class and the CAN overlay network generated is seen in Figure 1. Continuing

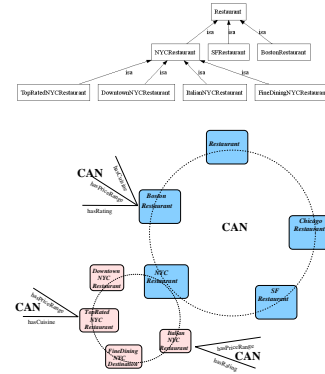


Figure 1: CAN overlay network

with the Restaurant example, as services register within CAN nodes and instances are created, they are classified into the subclasses of *Restaurant*. When a new CAN node joins the network, one of the CAN dimensions is split into two and data is transferred over to the new node. If there are  $c$  classes and  $d$  dimensions, classes are separated into  $d$  parts where each part contains  $c/d$  classes. According to some criteria, one of these dimensions is chosen and split into two. The current criteria we use is choosing the dimension with the largest number of keys. However, in the future we will implement network management techniques which keep track of the overloaded servers and split the dimension which has the greatest number of overloaded nodes. Thus, if the initial node has 3 dimensions with 10 classes in each dimension, then the range of each dimension is:  $[0 - 9]$ ,  $[0 - 9]$ ,  $[0 - 9]$ . When a new node joins the network, one of the dimensions is split and the resulting two nodes will have the following range of values:  $[0 - 4]$ ,  $[0 - 9]$ ,  $[0 - 9]$  and  $[5 - 9]$ ,  $[0 - 9]$ ,  $[0 - 9]$ .

Establishing a global service discovery architecture such as GloServ allows us to use ontologies to build a scalable, logically connected network. Ontology queries are then issued in GloServ and are efficiently routed to the correct server. Below we describe the ontology querying mechanism and the recent enhancement of combining the ontology queries with keyword search.

## 5. ONTOLOGY QUERYING

A service registration instance is distributed to all CAN nodes that handle the service classes it belongs to. Since we are using a CAN distributed hash table, not every node within the system needs to be updated. For queries, when a query is matched exactly, the first matching node will have the complete data set for that particular query restriction and thus further nodes need not be traversed. For a related match query, only the servers that hold logically similar information will be searched. Figure 2 gives a graphical overview of the query propagation in the CAN. We explain the details of ontology querying below by looking at the *Restaurant* ontology as our running example.




Figure 2: Query propagation in the GloServ CAN

## 5.1 Querying with Restricted Class Dimensions

A user initially contacts a GloServ user agent and enters a service name. The initial GloServer is found after following the steps outlined in Figure ?? . Since each hierarchical node handles a class which is disjoint from its siblings, the query is routed down only one branch reducing the query hops considerably. Once the correct GloServer is contacted, the user agent obtains the ontology pertaining to that service class. The interface to the user can either be human-centric or automated, depending on the implementation. In either case, a query is formed and sent to the GloServer. The query is a first order predicate logic statement that contains restrictions on various properties such as:

*(hasLocation some NYC)* and *(hasCuisine some (Korean or Chinese))*

The restaurant server creates a class with this query restriction and classifies it in its ontology. Since the subclasses of the *Restaurant* class are restricted by location, the query class gets classified as a subclass of the *NYCRestaurant* class. The query is then forwarded to the nodes that handle *NYCRestaurant* classes. When a node is found, the query class is classified again. Since the *NYCRestaurant* class has subclasses that have cuisine restrictions, when the reasoner classifies the query class, it becomes a subclass of *NYCRestaurant* and a superclass of *KoreanNYCRestaurant* and *ChineseNYCRestaurant* classes. The classification indicates that the query must be routed to the servers handling Chinese and Korean restaurants. In order to route the query within a CAN, the query needs to reduce to a dimension and key. We use the dimension and key values assigned to each of these classes during the CAN network generation to convert the ontology class to a  $\langle dimension, key \rangle$  pair

We illustrate ontology querying with the following example. Let us assume we have a *NYCRestaurant* ontology that has 30 subclasses, separated into 3 dimensions with 10 subclasses in each dimension. Furthermore, the *ChineseNYCRestaurant* subclass is assigned to dimension 0 with key 0 and *KoreanNYCRestaurant* to dimension 1 with key 0. If a user queries for:

*(hasLocation some NYC)* and *(hasCuisine some (Korean or Chinese))*

the query message is  $[0; 0; *]$ . As seen in Figure 2, Node1 receives the query and stops propagating it because it handles these classes. If a user relaxes her query requirements to not only include equivalent, superclass or subclass relationships but sibling relationships as well, Node1 looks at the sibling classes and issues query messages for each. For example, if the query message  $[*; 4; *]$  comes

into Node1 where semantic matches to the query are classes that are numbered 4, 5 and 6 in dimension 1, then the query message is converted to  $[*; 4, 5, 6; *]$ , processed in Node1 and propagated to Node3. A query continues to propagate until the original node is reached. Since a dimension is circular, it is guaranteed that the query will return back to its original position with at most  $O(n^{1/d})$  hops.

## 5.2 Query Matching of Related Information

GloServers receive queries in first order predicate logic statements. Query languages such as RQL [17] or SPARQL [4] are not used because our query matching algorithm does not just look for exact matches. Rather, as mentioned above, the GloServer creates a *query* class with the logical restriction specified in the query and classifies this *query* class within the ontology. The *query* class's superclasses, equivalent classes and subclasses are analyzed. For exact query matching, the equivalent classes and the subclasses are looked into. For related query matching, the superclass's children (which are the restricted class's siblings) are analyzed. Each of these siblings have certain restrictions on various properties. The related query matching algorithm finds properties that are related to the *query* class's properties and looks into the siblings that have these property restrictions.

Each property has a domain class and a range class. In order to find a related property, the range is classified and the equivalent classes and subclasses of the range are looked into. For example, the *Cuisine* class has the subclass *Italian* which has subclasses *Pizza* and *Pasta*. When a query comes in for a pizzeria with a five star rating in NYC, the query class will have the following restriction:

*(hasLocation some NYC)* and *(hasCuisine some Pizza)* and *(hasRating some FiveStar)*

This query class is classified according to how the ontology is constructed. In our ontology, it first gets classified under the *ItalianNYCRestaurant* class. If there are no instances within this class that have a *Pizza* cuisine and *FiveStar* rating, then the related classes of the the *Pizza* class are analyzed. Since the *Pasta* class is related to the *Pizza* class, the query is reformulated to include *Pasta* as the cuisine. Alternately, a user may choose to have related information in the query even if exact ones exist in which case the results given are both exact and related matches.

## 5.3 Querying a CAN with Property Dimensions

In the previous example, we looked at queries that were mapped to classes which had restricted subclasses mapped to a CAN. As the class restriction narrows, it may not be necessary to further restrict classes. But as the registration and query load grows within these servers, it is best to distribute the data where each dimension is a property type. For this case, querying is a bit different. Since we do not have subclasses to classify the query class by, we must look at the query class itself and generate keys to distribute within the CAN.

From the previous example, the query class lands in the nodes that contain the *ChineseNYCRestaurant* and *KoreanNYCRestaurant* classes. If these classes are not broken down further into subclasses, then the remaining unrestricted properties are *hasRating* and *hasPriceRange* properties. Thus, a 2-dimensional CAN is generated where each dimension represents a property. If the *hasRating* property has five values,  $[OneStar, TwoStar, ThreeStar, FourStar, FiveStar]$ , and *hasPriceRange* has four values  $[InExpensive, Mod-$

erate, Expensive, VeryExpensive], then there are a total of  $5 \times 4 = 20$  possible query combinations to issue. If the query was more specific, where a price range was specified, then the `hasPriceRange` property value is fixed and only five queries are issued.

Once the query is routed to the correct nodes, it is classified and all the inferred instances are obtained which match this query. The instances are then further analyzed by doing key word matching as we describe below.

## 6. KEY WORD MATCHING

Thus far, we have described ontology-based queries for query propagation and matching of service instances. However, services may also want to describe themselves with free text, such as with key words that are not already defined in the ontology. In order to be able to handle this case there needs to be a way for the ontology to handle key words and concepts. Below we describe an algorithm on how to incorporate key words in service registrations and queries.

### 6.1 Service Registration with Key Words

When a service registers within a given service class in `GloServ`, an instance is created, properties are populated and the instance is classified under a number of restricted classes in the service ontology. Restricted classes are normally restricted by object properties because the range of these properties are predetermined classes. The service instance is then routed to the servers which hold information on these restricted classes. We discuss service registration with key words for object properties and continue to look at an example using the `Restaurant` class.

For the `Restaurant` ontology, we restrict the classes by the `Neighborhood` and `Cuisine` classes. Thus, if this service provider is a Chinese restaurant in NYC, it is classified under a class which has as its restriction:

(*hasNeighborhood some NYC*) and (*hasCuisine some Chinese*)

The `hasCuisine` property has its range set to the `Cuisine` class and the `hasNeighborhood` property has its range set to the `Neighborhood` class. These classes can then have a set of nested subclasses. For example, the `Cuisine` class can have the subclass `Asian` which can then have the subclasses `Chinese`, `Japanese` and `Korean`. Although the `Cuisine` class can be constructed to be very rich in its subclass definitions, this is not always guaranteed. Thus, we would like to allow services to provide extra information when registering to include specific key words. For example, when a Chinese restaurant is registering, it sets its `hasCuisine` property to `Chinese` and then should have the option of adding extra keywords which may include their most popular menu items, daily specials, etc. Although all object type properties can have key words added to them, this may not be necessary. For example, the `Restaurant` ontology has a number of object properties such as `hasNeighborhood`, `hasCuisine`, `hasRating`, `hasPriceRange`. The `hasRating` or `hasNeighborhood` properties do not vary much in terms of expressivity. Thus it is redundant to add key words to these properties. Since OWL DL allows properties to be annotated with properties themselves, if a property is annotated with a `keyWordMatch=true` annotation, then key word matches are allowed for those properties.

In order to accomplish this, we create a `KeyWord` class which holds a list of key words that services have created while registering. The service provider fills out the object properties and is then given the option of creating key words for each of these properties. Thus, if the service provider sets the `hasCuisine` restaurant value

to `Chinese`, it is prompted with a list of key words from the `KeyWord` class which it can choose from and tag onto the `hasCuisine` property. It is also given the choice of creating new key words, which are added as new terms within the `KeyWord` class.

The property `hasKeyWord` is an object property which points to the `KeyWord` class. Every class in the ontology, which can be tagged with key words, has the `hasKeyWord` property as one of its properties. When a service chooses the `Chinese` class as its cuisine, an instance of the `Chinese` class is assigned to that service's `hasCuisine` property. If the service provider wants to add extra key words describing specifics to the Chinese cuisine, it is given a list of already generated key words it can choose from. The service provider chooses from this list and adds these key words to the instance of the `Chinese` cuisine class by inserting multiple `hasKeyWord` properties to the instance. If the service provider wants to add new key words, these get added to the `KeyWord` class and are tagged onto the `Chinese` cuisine instance.

Additionally, the service provider can add a set of key words which are not tied to properties but generically describe its own service. In this case it will populate the `hasKeyWord` property for its own service instance.

### 6.2 Querying for Services with Key Words

A query which includes key words is constructed in two parts. The first part is the ontology-based query. The second part holds key word information for each property. This is then converted to a complete ontology query.

If the front-end to the system is a graphical user interface, a user fills out the ontology form and can then enter key words next to each of the properties which has key word matching enabled. This query is sent to the back-end in two parts: ontology-based query and key word query. The ontology query is classified and sent to the appropriate `GloServers`. A new ontology query is then constructed to also include the key word queries.

This ontology query is constructed by first matching the key words to the key words in the `KeyWord` class. This can either be done by asking the user to choose a list of key words directly from the `KeyWord` class or have the user enter random key words. For the first case, the ontology query must be issued first in order for it to be routed to the appropriate server. Then a list of terms from the `KeyWord` class within that server is returned to the user so that she can choose those that match her request. For the second case, a user can add the key word terms along with the ontology query because the key words are matched to the terms in the `KeyWord` class using a text matching tool. Our implementation handles the first case, but can easily be extended to handle the second one by adding a text matching engine to match the key words to the key words already in the `KeyWord` class. This is a desirable feature because it extends to different types of front-ends which may not be user centric.

Once the key words are determined, an ontology query is built for each of the properties. For example, let us say the user entered the key words `Szechuan` and `Cantonese` as additional key words for the `hasCuisine` property. A restricted subclass is formed under the `Cuisine/Chinese` class with the restriction:

(*hasKeyWord has Szechuan*) and (*hasKeyWord has Cantonese*)

The ontology is classified and a list of inferred instances of the `Chinese` class which have these key words are classified under the query class. The name of an instance is usually the name of its class with a unique number to distinguish it from other instances in that class. Thus, a list of instances returned could be: `Chinese_1` and



Chinese\_2. Once these instances are obtained, the original ontology query is changed to include these specific instances. Thus, the original query:

(*hasCuisine* **some** Chinese) is replaced with:  
 (*hasCuisine* **has** Chinese\_1) or (*hasCuisine* **has** Chinese\_2)

A restricted class is created under the *Restaurant* class with this condition and the reasoner is run on the ontology to obtain a list of inferred instances which match this restriction. These instances are returned to the user. Additionally, since the key word queries match the *Chinese* class, related query matching can also be done as described above.

Besides entering key words for specific properties, the user may enter key words which give generic descriptions of the service. For these generic key words, the original ontology query is extended to include a condition for the *hasKeyWord* property. If the user enters key words such as: *rotating* and *view*, the ontology query example above is changed to:

(*hasCuisine* **some** Chinese) and  
 ((*hasKeyWord* **has** rotating) or (*hasKeyWord* **has** view))

A query class is created under the *Restaurant* class with this restriction and the ontology is classified to obtain all the instances which have these key words.

If users enter key words that do not match key words already in the ontology, then the instances which match the ontology part of the query are returned to the user and the key words that do not have a match are discarded.

### 6.3 Discussion

As discussed in Section 3, we try to limit the size of each service ontology so that classification will be done in a reasonable amount of time. When adding key words to an ontology, this can potentially cause the ontology to become very large. Below we discuss how to solve this problem.

First, periodic text matching is done between key words so that similar key words are merged into one term. Second, key words that are queried for often are included in the ontology and those that do not appear as often are stored only in the database instances. Thus, storing key words in both the ontology directly and the database, gives us the benefit of doing reasoning with key words, but also limiting the size of the ontology. Third, we store the *KeyWord* class apart from the main service ontology and merge it with the service ontology only when a key word search is being performed. Implementing these optimizations are part of our future work.

## 7. IMPLEMENTATION

We have implemented a prototype of GloServ using Protege [11] and Racer-Pro [15]. Protege is an open-source development environment for ontologies and knowledge-based systems. In order to follow a real-world classification, we have written tools to automatically generate ontologies pertaining to the restaurant classification in <http://www.menupages.com>. The *Restaurant* ontology is modified to represent the CAN lookup table. The subclasses within *Restaurant* are assigned to a unique  $\langle dimension, key \rangle$  pair. When a node joins a server, the server's ontology is split across a dimension and transferred over to the new node.

As the CAN is generated, nodes enter the system and are assigned to a zone. Each server initially holds many classes but as the number of nodes increase, the servers hold one class per dimension. Once a class exceeds a threshold for registration or querying,

it checks with other servers that handle the same class to see if a CAN subnetwork has already been formed. If it has, it caches the subnetwork's supernode information and transfers its data to this subnetwork. Subsequent registrations and queries are sent to this subnetwork. Otherwise, if there is no subnetwork, it processes its preconfigured ontology to generate the CAN subnetwork, and transfers data there. When subclasses do not exist, it parses the unrestricted properties and generates a CAN with property dimensions. We have built an ontology generating tool which automatically generates an ontology given a configuration file. A full CAN network is generated automatically when given the following input: an ontology configuration file, number of nodes in the CAN, list of server host names or IP addresses, list of super nodes in the CAN. The CAN is generated by the main server invoking remote "JOIN" messages from the list of listening servers.

We have built a front-end user interface to GloServ written in PHP. The interface has Google Maps which displays location-based services. The front-end parses the OWL ontology using PHP and automatically generates a form based on the service description given in the ontology. Services providers and users register or query via this form. The ontology also has annotation properties which guides the front-end server in populating the fields within the ontology. Thus, for properties that may also have key word search enabled, the interface displays a text box underneath the form for the user to enter key words in. The interface is seen in Figure 3 and can be accessed on the web at <http://glo.serv.dyndns.org:8080>

Tests for the ontology querying have produced promising results. For an ontology with 100 classes, the time it took to classify ranged from 0 to 15 ms. Query propagation across the CAN for 10 nodes averaged 80 ms. We estimate the CAN networks to have approximately 1000 nodes each since they represent a subdomain of the full ontology, following from the location-based example, a CAN network for *NYCRestaurant* classes. Thus, the maximum delay would average 8 seconds. The full query including reading 10,000 instances in the database took around 150 ms, using MySQL as the back-end database. We are still in the process of testing the key



Figure 3: GloServ Front-End

word searching extension and will test the scalability of the queries by running experiments to examine the latency of the query routing. We will then compare results from pure ontology queries to those that have been enhanced with text searches to see if these yield more accurate results.

## 8. RELATED WORK

### 8.1 Service Discovery

Current approaches to service discovery are either centralized or decentralized. Centralized service discovery systems include: SLP [14], Jini [20], and UDDI [5]. These systems are limited in network scalability as well as in representing service data.

The existing work closest to our research use schemas to map onto a network [19], [7], [26]. These systems are similar to GloServ in that they use data from a schema or ontology to map onto a network. [7] outlines a semantic gossiping framework that exchanges ontology information within a peer-to-peer network. It uses Gnutella [12] as its underlying peer-to-peer structure. The main problem with this system is that it uses flooding to broadcast its queries and thus reduces the scalability of the network.

[19] proposes the HyperCuP, a 2-tier peer-to-peer hierarchy which uses indices within a super-peer topology. The indices are built using schema information from associated peers. Super-peers are connected to each other in a hypercube or cayley graph and represent disjoint concepts. Underlying peers of a super-peer contain information regarding that particular concept. [26] describes Meteor-S which is also similar in that it separates peer-to-peer networks into different domains and represents these services using an ontology. However, the JXTA framework's Peer Discovery Protocol is used which employs flooding techniques for disseminating information and hence does not scale globally.

GloServ differs from these systems by using an ontology to map a multi-tier hierarchical peer-to-peer network. It is similar to [19] and [26] in that a concept represents a sub-network. However, unlike [19], concepts that are disjoint from each other are hierarchically organized whereas concepts that are similar to each other are organized in a CAN. Indices within the CAN are formulated according to the ontological content of each node whereas in [19], indices refer to whole peers. Also, because of the CAN DHT, GloServ generates  $O(\log_2 n)$  query messages whereas in [26] flooding is used which generates  $O(n)$  messages which is inefficient on a global scale. Additionally, [7] and [19] describe the peer-to-peer network formation, namely, dealing with nodes entering and leaving the system and concentrate less on describing how the data is disseminated. We, on the other hand, describe the formation of the network as well as describe algorithms that distribute and query the data by mapping the ontology onto a CAN network.

## 8.2 Ontology-based Information Retrieval

Currently, work done in combining ontology-based search with information retrieval focuses on adding semantic meaning to documents. The key words are already defined within these documents and they are then mapped into the ontology and classified within certain domains. A few of these systems, among many others, are described in [21], [6] and [18].

GloServ addresses a different problem. It seeks to represent and discover services using ontology queries and key word search. Service data is already represented as ontology instances. Thus, a set of key words does not exist initially, but the ontology is modified as services register and add their own key words to the ontology. Since the service classification ontology is used to distribute data in peer-to-peer overlay networks, the set of key words generated, as services register, will belong to a certain number of service classes handled by that server. Thus, key word generation and search is dynamic and can apply to all service domains.

## 9. CONCLUSION

GloServ is a hierarchical peer-to-peer global service discovery system using OWL DL. GloServ functions both on a wide area as well as a local area network. Broad range of services are defined flexibly using OWL ontologies. The GloServ architecture achieves

large-scale distribution of semantic data that is queried for with specificity and efficiency. The ability to reason in OWL DL promotes intelligent distribution of service content across nodes connected in a CAN peer-to-peer network.

We have described a recent enhancement to GloServ which combines ontology querying with key word search. The ontology query is used to route the query to the servers that hold information on these services and to find a list of matching and related instances. Key words are then matched to the properties of these instances. Combining ontology querying with text searching enhances the description and discovery of services.

## 10. ACKNOWLEDGEMENT

We would like to thank Peter F. Patel-Schneider of Bell Labs Research for his contribution to this work.

## 11. REFERENCES

- [1] 2000 census: US municipalities over 50,000: Alphabetical. <http://www.demographia.com/db-2000city50k.htm>.
- [2] Demographia world urban areas. <http://www.demographia.com/db-worldua.pdf>.
- [3] Owl web ontology language. OWL <http://www.w3.org/2004/OWL/>.
- [4] SPARQL query language for rdf. <http://www.w3.org/TR/rdf-sparql-query/>.
- [5] UDDI technical white paper. white paper, UDDI (universal description, discovery and integration), September 2000. <http://www.uddi.org/pubs/>.
- [6] Jose Maria Abasolo and Mario Gomez. MELISA: An ontology-based agent for information retrieval in medicine. *Proceedings of ECDL 2000 Workshop on the Semantic Web*, September 2000.
- [7] K. Aberer, P. Cudre-Mauroux, and M. Hauswirth. A framework for semantic gossiping. In *ACM SIGMOD: Special Interest Group on Management of Data*, 2002.
- [8] Knarig Arabshian and Henning Schulzrinne. Distributed context-aware agent architecture for global service discovery. *The Second International Workshop on Semantic Web Technology For Ubiquitous and Mobile Applications (SWUMA'06)*, 2006.
- [9] Knarig Arabshian and Henning Schulzrinne. An ontology-based hierarchical Peer-to-Peer global service discovery system. *Journal of Ubiquitous Computing and Intelligence (JUCI)*, 2006.
- [10] Tom Gardiner, Ian Horrocks, and Dmitry Tsarkov. Automated benchmarking of description logic reasoners. In *Proc. of the 2006 Description Logic Workshop (DL 2006)*, volume 189, 2006.
- [11] J. Gennari, Mark A. Musen, R. W. Ferguson, W. E. Grosso, M. Crubézy, H. Eriksson, N. F. Noy, and S.-C. Tu. Evolution of protégé: An environment for knowledge-based systems development. Technical report, Stanford University, 2002.
- [12] Gnutella. <http://gnutella.wego.com>.
- [13] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
- [14] E. Guttman, C. Perkins, J. Veizades, and M. Day. Service location protocol, version 2. RFC 2608, Internet Engineering Task Force, June 1999.
- [15] Volker Haarslev and Ralph Moller. Racer user's guide and reference manual version 1.7.19. Technical report, Technical

University of Hamburg-Harburg, University of Hamburg, 2004.

- [16] Matthew Horridge, Alan Rector, Nick Drummond, Holger Knublauch, and Hai Wang. A user oriented owl development environment designed to implement common patterns and minimise common errors. In *3rd International Semantic Web Conference (ISWC2004)*, Hiroshima, Japan, Nov 2004.
- [17] Gregory Karvounarakis, Sofia Alexaki, Vassilis Christophides, Dimitris Plexousakis, and Michel Scholl. RQL: a declarative query language for RDF. In *Proceedings of the 11th International World Wide Web Conference*, pages 592–603, 2002.
- [18] Latifur Khan, Dennis McLeod, and Eduard Hovy. Retrieval effectiveness of an ontology-based model for information selection. *The VLDB Journal*, 13(1):71–85, 2004.
- [19] Alexander Loser, Wolf Siberski, Martin Wolpers, and Wolfgang Nejdl. Information integration in schema-based peer-to-peer networks. In *Proceedings of the Conference on Advanced Information Systems Engineering*, June 2003.
- [20] Sun Microsystems. Jini architectural overview. Technical report, 1999.
- [21] Hans-Michael Muller, Eimear E. Kenny, and Paul W. Sternberg. Textpresso: An ontology-based information retrieval and extraction system for biological literature. *PLoS Biology*, 2, 2004.
- [22] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, and Scott Shenker. A scalable content-addressable network. San Diego, CA, USA, August 2001. ACM.
- [23] Alan Rector. Modularisation of domain ontologies implemented in description logics and related formalisms including owl. In *2nd International Conference on Knowledge Capture (K-CAP)*, Sanibel Island, FL, 2003.
- [24] Antony Rowstron and Peter Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pages 329–350, Heidelberg, Germany, November 2001.
- [25] Ion Stoica, Robert Morris, David R. Karger, Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet applications. San Diego, CA, USA, August 2001. ACM.
- [26] Kunal Verma, Kaarthik Sivashanmugam, Amit Sheth, Abhijit Patil, Swapna Oundhakar, and John Miller. METEOR-S WSDI: A scalable p2p infrastructure of registries for semantic publication and discovery of web services. *Journal of Information Technology and Management. Special Issue on Universal Global Integration*, 6(1):17–39, 2005.